

recognize that there are many possible hierarchies within the scope of the present invention. For example, in systems without RAID volume controllers, the VLUN 32 may write and read data directly to the controllers for the disk drives 36 and 38. Similarly the requests to the VLUN 32 may originate with a database instead of a file system 30.

The VLUN 4 is responsible for determining an available storage location for data sent from a user, such as a file system 1 or database, of the host electronic device 2. **Figure 3** depicts the sequence of steps utilized by an illustrative embodiment of the present invention to store and then read data using a VLUN 4. A user on a host electronic device 2 sends data to the VLUN 4 for storage (step 42). The VLUN 4 determines available storage locations and then sends the data to a RAID volume controller (step 44) for an available RAID set. The RAID volume controller sends the data to a RAID set (step 46). Subsequently the VLUN 4 receives a request for data access from the user (step 48). The VLUN 4 retrieves the data from the RAID volume controller and forwards it to the requesting user (step 50). Those skilled in the art will recognize that alternate implementations of the present invention are possible. For example, the VLUN 4 may send the data to more than one RAID volume controller for storage, thus implementing different RAID levels for the same data.

The VLUN 4 acts to ensure continued data access for a user in the event operations to reconfigure or repair the storage mediums are required or advisable. **Figure 4** depicts the sequence of events followed by an illustrative embodiment of the present invention upon the detection of an error in one side of a RAID set. A user on a host electronic device 2 sends data to the VLUN 4 for storage (step 54). The VLUN 4 sends the data to a selected RAID volume controller (step 56). The RAID volume controller sends the data to a RAID mirror set where multiple copies of the data are stored (step 58). Depending upon the RAID level of the RAID set, the copies may be two complete copies of the data, or may be a complete copy and a parity copy. The copies may or may not be striped across the storage mediums. In some embodiments, two disk drives may be used for the RAID mirror set. In other embodiments three or more whole disk drives may be used to create a RAID set. Alternatively, segments of disk drives with adjacent logical memory addresses may be utilized by a RAID volume controller to create a software RAID set. Those skilled in the art will recognize that there are a multitude of possible components, and many different configurations of

components, that may be utilized without departing from the scope of the current invention.

As an illustrative example, subsequently to storing the data, the RAID controller
5 may detect an error in one spindle of the RAID set. The VLUN 4 is notified of the error
in the RAID set (step 60). Once the error in the spindle of the RAID set has been
detected, the method of the present invention provides multiple ways to recover from the
detected error. In one embodiment, the RAID set is repaired with a spare spindle (i.e.,
spare disk drive) (step 62). If the error is detected during an access request from the
10 VLUN 4, a valid copy of the data is sent to the VLUN before or during the repairs to
the spindle with the error. When the new spindle is in place, the data from the valid
spindle of the RAID set is copied to the new spindle. In another embodiment of the
present invention, data is retrieved from the valid spindle of the RAID set via the RAID
volume controller and sent to the VLUN 4. The VLUN 4 verifies available storage
15 locations and sends the data to a new RAID volume controller to be placed into a new
RAID set (step 64). In another embodiment, data is retrieved by the VLUN 4 from the
valid spindle of the RAID set using the RAID volume controller. Data is thereafter sent
to a controller for a software RAID set which has been created using multiple storage
mediums (step 66). Similarly, the VLUN 4 may schedule the movement of data from
20 one RAID set to another so as to balance hardware use. Requests for data received
during scheduled reconfiguration are treated as in the error detection process described
above. Those skilled in the art will recognize that there are multiple error recovery and
reconfiguration procedures possible within the scope of the present invention and those
listed here are illustrative but not exhaustive.

25

In an alternative embodiment, the VLUN 4 sends data directly to available
storage mediums. **Figure 5** depicts the sequence of steps followed to allocate data
directly from the VLUN 4. The VLUN 4 receives data from a user for storage
(step 70). The VLUN 4 checks the available storage mediums, such as disk drives,
30 both local and remote (step 72) for available storage space. The VLUN 4 sends the
data directly to the controllers for the storage mediums (step 74). There are a number
of ways of storing the data within the scope of the current invention. The data may be
striped across multiple disks (step 76) or the data may be mirrored among several disks
(step 78). Those skilled in the art will recognize that the data may be striped and
35 mirrored across multiple disks. Alternatively, complete duplicate copies of data or

parity data may be stored. In this embodiment, the VLUN 4 allocates the data directly instead of sending it to an intermediate controller, such as a RAID volume controller.

If the VLUN 4 is allocating data directly to the storage mediums without using a RAID volume controller, the VLUN is responsible for performing reconfiguration and repair operations on the storage mediums in a manner that minimizes data access disruptions to a user of the host electronic device 2. **Figure 6** depicts a sequence of steps followed by the VLUN 4 in an illustrative embodiment of the present invention upon detection of an error in a storage medium holding data directly allocated by the VLUN. The VLUN 4 allocates data directly to the storage medium in the manner discussed above (step 84). The VLUN then receives a user request to either read or write the data and attempts access to the stored data. Upon attempting access, the VLUN 4 detects an error in part of the stored data (step 86). The error may be mechanical error in the storage medium preventing the writing or reading of data, or the error may be a data error detected by the VLUN in comparing multiple copies of the data retrieved from the storage mediums which interferes with the reading of data. The VLUN 4 retrieves a valid copy of the data and sends it to the user (step 88). The VLUN 4 attempts the repair/reconfiguration of the detected error contemporaneously with retrieving the valid copy of the data or immediately thereafter. If the VLUN 4 is able to identify an additional available storage medium, the valid data may be copied to the identified storage medium (step 90). If the VLUN 4 is unable to identify a single available storage medium suitable for the valid data, the VLUN may segment the data and distributes the segments to multiple disks (step 92).

The illustrative embodiments of the present invention provide a virtual interface for the reading and writing of storage requests of a user of a host electronic device. By making the storage process transparent to the user, real-time reconfiguration and repair of storage mediums can take place without unduly disrupting user access to data. By providing the virtual interface on the host electronic device rather than locating the virtual interface out on the network, the user maintains greater access to data and more flexibility in recovering from hardware and software errors.

It will thus be seen that the invention attains the objects made apparent from the preceding description. Since certain changes may be made without departing from the scope of the present invention, it is intended that all matter contained in the above description or shown in the accompanying drawings be interpreted as illustrative and not